

HP InfiniBand solution for Oracle Application Server



Overview	2
Oracle Application Server overview	2
Advantages of InfiniBand in an Oracle Application Server environment	2
HP InfiniBand solution for Oracle Application Server	3
Solution overview	3
Sample HP InfiniBand configuration for Oracle AS 10G and Oracle RDBMS 10G.....	4
Installation process.....	5
Oracle AS on InfiniBand installation process.....	5
Hardware setup and configuration	5
Red Hat Enterprise Linux® installation	5
Installation of InfiniBand software components	5
Oracle RDBMS 10G installation.....	6
Oracle Application Server 10G installation.....	6
Transparent SDP configuration.....	7
Sample J2EE application installation	8
For more information.....	9

Overview

This whitepaper explains how to implement an Oracle Applications Server (AS) J2EE solution for the database connection using HP ProLiant servers, HP StorageWorks, and HP InfiniBand™ options. The whitepaper is intended for system administrators, system architects, and systems integrators who are considering the advantages of HP InfiniBand-based systems in an Oracle AS J2EE environment.

Oracle Application Server overview

Oracle AS is a fully integrated application server – including a J2EE Compatible container, an enterprise portal, wireless, business intelligence, single sign-on using a LDAP compliant directory, and much more. It is the integrated nature of Oracle AS that distinguishes it in the marketplace. Oracle strives not only to provide a container for J2EE applications, but to enable J2EE applications to be managed more easily, deployed as web services, accessed via a single sign-on, made more secure, and tied more closely to business level activities.

Typically, the Oracle AS is installed on its own server or set of servers and these servers access an Oracle database instance on another server across a network. The most typical network connection between Oracle AS server and the Oracle database server is via TCP/IP over Ethernet.

A typical Oracle Application Server system has the following components:

- One or more application servers
- Oracle database instance
- Public local area network (LAN)

Advantages of InfiniBand in an Oracle Application Server environment

The most common method of communication between application server machines and database server machines today is TCP/IP over Ethernet. TCP/IP and Ethernet both have limitations. The TCP/IP processing on the application server and database server machines can result in substantial CPU utilization on both machines with typical latencies of 50 to 100 microseconds (μ s). Additionally, observed Gigabit Ethernet throughput is limited to approximately 0.8 gigabits per second (Gbps) on a single network interface card (NIC).

InfiniBand, a standards-based alternative to Ethernet, overcomes the performance limitations of TCP/IP-based solutions. The key advantages of InfiniBand over TCP/IP and Ethernet is InfiniBand's high throughput, low CPU utilization, and ultra-low latency. The current InfiniBand bandwidth of 10 Gbps¹ provides observed performance over 8 Gbps. CPU utilization of only 1-3 % with latency less than 5 μ s is possible.

InfiniBand supports its own InfiniBand-specific protocol, sockets direct protocol (SDP). This is the standard wire protocol for the InfiniBand Architecture (IBA) for support of stream sockets (SOCK_STREAM) networking over IBA.

There are two different versions of SDP protocol that are supported by Oracle. The first is native SDP support, where the Oracle product implements the SDP directly. The second protocol, called Transparent SDP, is where the Oracle product is configured to use TCP, and the SDP protocol provider converts TCP connections to SDP without any configuration changes at the Oracle level. Native SDP provides support for asynchronous I/O, whereas Transparent SDP only supports synchronous I/O.

Because it supports asynchronous I/O, native SDP supports one of the most important features of InfiniBand – remote direct memory access (RDMA). RDMA is a communications technique that allows data to be transmitted from the memory of one computer to the memory of another computer without

¹ InfiniBand 10Gbps bandwidth is represented as 4x or four lanes of traffic each at 2.5 Gbps.

passing through either device's CPU, without needing extensive buffering, and without calling to an operating system kernel. RDMA can bring about a substantial reduction in CPU utilization on both the application and database servers.

Oracle's 10G RDBMS product supports native SDP; however, Oracle products that use an earlier version of the Oracle DB client can only utilize Transparent SDP. Both Oracle AS 9i and 10G Release 1 use the Oracle 9i DB client. Therefore, they can only use Transparent SDP. While this means that Oracle AS cannot take advantage of asynchronous I/O or RDMA today, it still benefits from the additional bandwidth and higher throughput provided by InfiniBand with Transparent SDP. Oracle AS 10G Release 2 will support native SDP.

Applications that connect to the database by using thin JDBC driver are also restricted to Transparent SDP. Applications that use the OCI driver can use native SDP protocol provided the product support SDP natively.

HP InfiniBand solution for Oracle Application Server

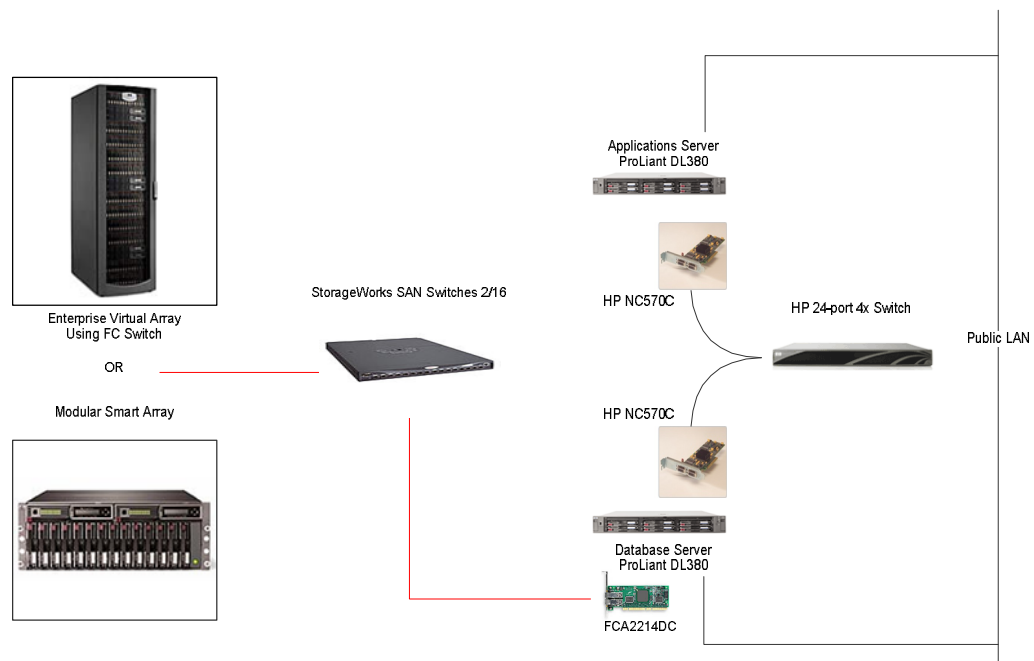
Solution overview

Many HP customers run J2EE applications in an Oracle Application Server environment using standards-based HP ProLiant servers and HP StorageWorks Fibre Channel connectivity products. With the addition of HP InfiniBand options for ProLiant servers, customers can implement a complete HP-based Oracle Application Server solution using InfiniBand for the database connection.

Figure 1 depicts a sample configuration composed of the following components:

- Application Server
 - HP ProLiant DL380 G4 server
 - (4) 36GB disk drives
- Database Server
 - HP ProLiant DL380 G4 server
 - (2) 36GB disk drives
 - HP StorageWorks solution
 - § HP StorageWorks Modular Smart Array (MSA) or Enterprise Virtual Array (EVA)
 - § HP StorageWorks FCA2214DC Fiber Channel HBA
 - § HP StorageWorks SAN Switch
- HP InfiniBand cluster interconnect
 - HP NC570C dual port PCI-X InfiniBand HCA
 - HP 24 port 4x Copper Fabric Switch

Figure 1. HP InfiniBand solution for Oracle AS high level diagram



Sample HP InfiniBand configuration for Oracle AS 10G and Oracle RDBMS 10G

The orderable components necessary to build a 2-node sample configuration based on the HP ProLiant DL380 server, HP StorageWorks MSA 1000 for database storage, and HP ProLiant InfiniBand options for the private interconnect is shown in Table 1.

Table 1. Sample configuration bill of materials

Quantity	Part Number	Description
HP ProLiant Server Nodes		
2	361011-001	DL380R04 X3.6-1MB/800, 2GB High Performance
2	311584-B21	Intel® Xeon™ X3.6-1MB/800MHz Processor Option Kit
4	343056-B21	2GB of Advanced PC2 PC3200 DDR2 SDRAM DIMM Memory Kit (2 x 1024 MB)
6	286776-B22	36.4GB 15,000 rpm, U320 Universal Hard Drive, 1" (internal server storage)
HP InfiniBand Private Interconnect		
2	376158-B21	HP NC570C PCI-X Dual-port 4x Fabric Adapter
2	376232-B21	HP 1m 4x Fabric Copper Cable
1	376227-B21	HP 24-Port 4x Fabric Copper Switch
HP StorageWorks Database Storage		

Quantity	Part Number	Description
1	321835-B21	StorageWorks 2 Gb, Dual Port, 64-Bit/133 MHz PCI-X-to-Fibre Channel Host Bus Adapter for W2K, Windows Server 2003 and Linux®
1	221692-B22	5 Meter FC LC-LC Multi-Mode cable
1	201723-B22	Modular SAN Array 1000 with 256 MB cache
1	190211-001	StorageWorks Enclosure Model 4354R – Rack
14	286776-B22	36.4GB 15,000 rpm, U320 Universal Hard Drive, 1" (shared storage)

Installation process

Oracle AS on InfiniBand installation process

The following is a list of the installation steps required to implement Oracle AS on InfiniBand using ProLiant servers.

Hardware setup and configuration

Follow the appropriate installation guides to prepare your site and rack mount the HP StorageWorks storage device, HP StorageWorks Fibre Channel switch, HP ProLiant servers, and HP 24 port 4x InfiniBand switch. See [HP StorageWorks - Entry-level Storage Products](http://h18006.www1.hp.com/storage/entrystorage.html), <http://h18006.www1.hp.com/storage/entrystorage.html>, for the StorageWorks MSA 1000 storage device and related switches. See [HP.com - HP ProLiant InfiniBand™](http://h18004.www1.hp.com/products/servers/networking/index-ib.html), <http://h18004.www1.hp.com/products/servers/networking/index-ib.html>, for the InfiniBand switch. See [HP.com - ProLiant servers - Industry standard servers](http://h18004.www1.hp.com/products/servers/platforms/), <http://h18004.www1.hp.com/products/servers/platforms/>, for information on ProLiant servers. Use the HP FC HBA installation guide to install the Fibre Channel HBA in the database server and connect the Fibre Channel cables between the storage device, Fibre Channel switch, and the Fibre Channel HBA. See [HP Storage SAN Infrastructure Products](http://h18006.www1.hp.com/storage/saninfrastructure.html), <http://h18006.www1.hp.com/storage/saninfrastructure.html>, for HBA. Use the HP InfiniBand 4x HCA install guide to install the HCA in each server and connect the InfiniBand cables to the InfiniBand switch and both HCAs. See, [HP.com - HP ProLiant InfiniBand™](http://h18004.www1.hp.com/products/servers/networking/index-ib.html), <http://h18004.www1.hp.com/products/servers/networking/index-ib.html> for InfiniBand.

Configure the HP StorageWorks storage device for use with Oracle 10G database using the HP StorageWorks configuration guide for the storage device you have chosen.

Red Hat Enterprise Linux® installation

To install Linux, follow the Red Hat Enterprise Linux installation guides for both application and database servers and apply any patches deemed necessary. HP also provides a Linux installation guide specifically designed for installing Linux on ProLiant servers at the following link.

<http://h200006.www2.hp.com/bc/docs/support/SupportManual/c00257687/c00257687.pdf>

Installation of InfiniBand software components

Follow the instructions in the Install and configure HP InfiniBand 4x HCA Drivers for Linux section of the HP 4x InfiniBand HCA User Guide. By default, the drivers will automatically install support for all of the protocols used by Oracle including SDP and Internet protocol over InfiniBand (IPoIB).

Once the HP InfiniBand 4x HCA drivers have been installed and configured on both servers, modify one of the unused Linux Ethernet configuration files (or make a copy) for use as the private interconnect. Edit the hosts file to include both "hostname_IB" and the IP address defined in the Linux Ethernet configuration files used for private interconnect.

Specify the private interconnect IP address and hostname_IB

Both servers

Modify the hosts file to include the node names and IP addresses used for the private interconnect.

```
# vi /etc/hosts
10.1.1.10    dbhostname_IB
10.1.1.15    appshostname_IB
```

Copy ifcfg-eth0 and rename it ifcfg-ib0

```
# cd /etc/sysconfig/network-scripts
# cp ifcfg-eth0 ifcfg-ib0
```

Database server

Modify ifcfg-ib0 to include only the items listed below.

```
# vi /etc/sysconfig/network-scripts/ifcfg-ib0
DEVICE=ib0
BOOTPROTO=static
IPADDR=10.1.1.10
NETMASK=255.255.255.0
ONBOOT=yes
```

Application server

Modify ifcfg-ib0 to include only the items listed below.

```
# vi /etc/sysconfig/network-scripts/ifcfg-ib0
DEVICE=ib0
BOOTPROTO=static
IPADDR=10.1.1.15
NETMASK=255.255.255.0
ONBOOT=yes
```

Oracle RDBMS 10G installation

Once the HP InfiniBand 4x private interconnect is installed, configured, and tested, Oracle10G can be installed using the procedure documented in the Oracle 10G installation guide, Installing Oracle Database 10G on Linux x86. The following is a link for Oracle latest installation guide.

http://www.oracle.com/technology/pub/articles/smiley_10qdb_install.html

Oracle Application Server 10G installation

Once the database server is operational, Oracle Application Server 10G can be installed using the procedures documented in the Oracle installation guide, Oracle Application Server 10G Installation Guide. The following is a link for Oracle latest installation guide.

http://download-west.oracle.com/docs/cd/B12428_13/install.904/install/toc.htm

Transparent SDP configuration

The following steps are very crucial and require absolute path statements; any typing errors could render the system inoperable. The preload file, `ld.so.preload` points to a library, `libsdp.so`, to be loaded during system startup. If the path is incorrect, library name is misspelled, or library does not exist a reinstall of the operating system may be required.

On the database server

1. Verify the existence of `/lib/libsdp.so`
2. Stop the database
3. Stop the listener
4. `# echo "/lib/libsdp.so" > /etc/ld.so.preload`
5. `# echo "match listen *:1521" >> /usr/local/topspin/etc/libsdp.conf`
6. `# echo "match destination *:1521" >> /usr/local/topspin/etc/libsdp.conf`
7. Edit the `listener.ora` file to include the private interconnect settings.

```
LISTENER =  
    (DESCRIPTION_LIST =  
        (DESCRIPTION =  
            (ADDRESS_LIST =  
                (ADDRESS = (PROTOCOL = SDP) (HOST = dbhostname_IB) (PORT = 1521))  
            )  
        )  
    )
```

`NET_ASYNC_IO=FALSE`

`SDP_ASYNC_IO=FALSE`

8. Edit the `sqlnet.ora` file to include the following lines.

`NET_ASYNC_IO=FALSE`

`SDP_ASYNC_IO=FALSE`

9. Start the database
10. Start the listener
11. Check `/proc/topspin/sdp/conn_main` (should have 3 entries)

On the application server

1. Verify the existence of `/lib/libsdp.so`
2. `# echo "/lib/libsdp.so" > /etc/ld.so.preload`
3. `# echo "match listen *:1521" >> /usr/local/topspin/etc/libsdp.conf`
4. `# echo "match destination *:1521" >> /usr/local/topspin/etc/libsdp.conf`
5. Edit the `tnsnames.ora` file to included the private interconnect settings.

```
INT =  
    (DESCRIPTOIN =  
        (ADDRESS = (PROTOCOL = TCP) (HOST = dbhostname_IB) (PORT = 1521))  
        (CONNECT_DATA =
```

```
(SERVICE_NAME = INT)
)
)
```

6. Check /proc/topspin/sdp/conn_main (should have 1 entries)
7. Verify sqlplus connection with the database.

Sample J2EE application installation

Installation notes and scripts are generally provided with the J2EE application software.

Configure the J2EE application to utilize the private InfiniBand interconnect.

Modify the data sources for the sample J2EE Application as in the following example:

IASDBDS

JDBC URL jdbc:oracle:thin:@dbhostname_IB:1521:INT

Test the connectivity

1. sqlplus <Oracle user name>/<Oracle password>@<Oracle Service Name>
2. Browse to J2EE Application login screen

For more information

For additional information, refer to the resources detailed below.

Resource description	Web address
HP ProLiant InfiniBand options	http://h18004.www1.hp.com/products/servers/networking/index-ib.html
HP StorageWorks Modular Smart Arrays (MSA)	http://h18006.www1.hp.com/storage/arraysystems.html
HP ProLiant servers	http://h18004.www1.hp.com/products/servers/platforms/
HP Parallel Database Clusters (PDC) for Oracle RAC	http://h18004.www1.hp.com/solutions/enterprise/highavailability/oracle/index.html
HP support documentation and software downloads	http://welcome.hp.com/country/us/en/support.html
Oracle Application Server	http://www.oracle.com/appserver/index.html

© 2005 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

InfiniBand is a trademark or registered trademark of the InfiniBand Trade Association.

Intel and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a U.S. registered trademark of Linus Torvalds.

Oracle is a registered US trademark of Oracle Corporation, Redwood City, California.